

Data-Efficient Quickest Outlying Sequence Detection in Sensor Networks

Taposh Banerjee and Venugopal V. Veeravalli

Department of Electrical & Computer Engineering,
University of Illinois at Urbana-Champaign, Urbana, USA

Email: banerje5,vvv@illinois.edu

Abstract—A sensor network is considered where at each sensor a sequence of random variables is observed. At each time step, a processed version of the observations is transmitted from the sensors to a common node called the fusion center. At some unknown point in time the distribution of observations at an unknown subset of the sensor nodes changes. The objective is to detect the outlying sequences as quickly as possible, subject to constraints on the false alarm rate, the cost of observations taken at each sensor, and the cost of communication between the sensors and the fusion center. Minimax formulations are proposed for the above problem and algorithms are proposed that are shown to be asymptotically optimal for the proposed formulations, as the false alarm rate goes to zero. It is also shown, via numerical studies, that the proposed algorithms perform significantly better than those based on fractional sampling, in which the classical algorithms from the literature are used and the constraint on the cost of observations is met by using the outcome of a sequence of biased coin tosses, independent of the observation process.

Index Terms—Quickest change detection, observation control, minimax, multi-channel systems, outlying sequence detection, asymptotic optimality.

I. INTRODUCTION

In many engineering applications a sensor network is often deployed to observe a phenomenon, and to make statistical inference in a distributed and collaborative manner in real time. An application of particular interest is the detection of the onset of an activity or occurrence of an event in/around the object/phenomenon being monitored. For example, it is of interest to detect the arrival of an animal/bird to its habitat or to detect a sudden increase in the stress/strain on the infrastructure being monitored like bridges or buildings. Other applications include the detection of a bioterrorist attack, the detection of the arrival of an intruder in a geographical area, etc. The detection of such an event in real time can either be a primary objective of the network, or it could be used as a trigger to activate more sophisticated and costly monitoring systems or sensors, e.g., a video monitoring system or a human inspection. Such detection problems can be modeled under the framework of quickest change detection.

Parts of this paper have been presented at ICASSP 2013. This research was supported in part by the National Science Foundation under grant CCF 08-30169, CCF 11-11342 and DMS 12-22498, through the University of Illinois at Urbana-Champaign. This research was also supported in part by the U.S. Defense Threat Reduction Agency through subcontract 147755 at the University of Illinois from prime award HDTRA1-10-1-0086.

The problem of quickest change detection (QCD) is well studied in the literature; see [1], [2], [3], and [4] for a review of QCD. In the classical QCD problem there is a single sequence of random variables or vectors. Before the change the random vectors have a particular distribution. At some point in time, called the change point, the distribution of the vectors changes. The objective is to detect this change in distribution as quickly as possible, i.e., with minimum possible delay, subject to a constraint on the false alarm rate; see [5], [6], [7], [8], [9], [10], [11] and [12]. Depending on the availability of the information on the distribution of the change point, the QCD problem is either studied in the Bayesian setting of [8], or in non-Bayesian or minimax settings of [10] and [11].

In this paper we are interested in the decentralized version of the problem first studied in [13], and further investigated, for example, in [14] and [15]; see [4] and the references therein.

In the decentralized QCD model there is a set of sensors and a central decision maker called the fusion center. At each sensor a sequence of random variables is observed over time, and at each time step, a processed version of the observations is transmitted from each sensor to the fusion center. At the change point the distribution of the observations at all the sensor nodes changes. The objective in the decentralized model is to find a technique to process the observations locally at each sensor, and to find a fusion technique to be applied at the fusion center, to detect the change as quickly as possible, subject to a constraint on the false alarm rate.

In modern sensor networks there is a cost associated with acquiring observations at each sensor. Also, there is a cost associated with the communication between the sensors and the fusion center. That is there is a cost associated with acquiring data in the system. Thus, the change has to be detected in a *data-efficient* way. Moreover, a prior on the change point is often not available. Also often the change only affects a subset of the sensor nodes, and the information on the affected subset and even its size may not be known *a priori*. In the QCD literature the latter problem is called a multi-channel QCD problem.

The QCD problem in a multi-channel setting is studied in [16], [17] and [18]. However, in these papers neither the cost of observations at the sensors, nor the cost of communication between the sensor nodes and the fusion center is taken into account. In the papers [13], [14], [15], and [18], the cost of communication is controlled by restricting the amount of information transmitted to either one bit or few bits, but the

constraint on the cost of communication is not part of the problem formulation itself. The QCD problem where the cost of communication is considered explicitly is studied in [19] and [20]. However in these papers, the cost of observations at the sensors is not taken into account. Also the problems are not considered in a multi-channel setting, i.e., in these papers the change affects all the sensors at the time of change.

In [21] and [22] we extended the classical QCD formulations studied in [8], [10] and [11] by putting an additional constraint on the cost of observations used in the detection process. We proposed problem formulations, for the Bayesian setting in [21] and for two minimax settings in [22], where the objective is to minimize some version of the average delay, subject to constraints on the false alarm rate and a version of the average number of observations taken before the change point. For the i.i.d. model we proposed two-threshold extensions of the classical single-threshold algorithms, and showed that they are asymptotically optimal for the proposed formulations. We also showed via simulations that the two-threshold algorithms we proposed provide a significant gain in performance as compared to the approach of fractional sampling, in which the constraint on the observation cost is met by skipping samples randomly.

In [23] we extended the results from [22] to sensor networks where the change affects all the sensors. However in the problem formulations in [23], the cost of communication between the sensors and the fusion center is not taken into account.

In this paper we extend the results from [22] to a sensor network where the change affects the distribution of observations at an unknown subset of sensors. We refer to this problem as the quickest outlying sequence detection problem. We propose extensions of the minimax problem formulations from [10] and [11] to sensor networks by introducing additional constraints on the cost of observations used at each sensors, and the cost of communication between the sensors and the fusion center. We propose two algorithms: the DE-Censor-Max algorithm and the DE-Censor-Sum algorithm. Both the algorithms are based on the DE-CuSum algorithm we proposed in [22]. The DE-CuSum algorithm can be used for data-efficient QCD in a single sequence of observations. In both the algorithms we propose in this paper, the DE-CuSum algorithm is used locally at each sensor; thus ensuring data-efficiency at the sensors. In both the algorithms the local DE-CuSum statistic is transmitted from the sensors to the fusion center, if the local DE-CuSum statistic is above a certain threshold; this is censoring. In the DE-Censor-Max algorithm a change is declared at the fusion center when the *maximum* of the transmitted DE-CuSum statistics across the streams is above a threshold. In the DE-Censor-Sum algorithm a change is declared at the fusion center when the *sum* of the transmitted DE-CuSum statistics across the streams is above a threshold. We will provide detailed performance analysis of these algorithms. The analysis will reveal that the DE-Censor-Max algorithm is asymptotically optimal for the problems proposed, when the change affects *exactly one* stream, as the false alarm rate goes to zero. Also, using the results in [17], the DE-Censor-Sum algorithm is uniformly asymptotically optimal, for each

possible post-change scenario, as the false alarm rate goes to zero. We will also provide numerical results to compare the performance of the two proposed algorithms as a function of the number of outlying streams.

II. CENTRALIZED MINIMAX FORMULATIONS FOR DE-QCD AND THE DE-CUSUM ALGORITHM

Since the formulations and algorithm proposed in this paper crucially depend on the formulations and the algorithm proposed in [22], in this section we provide a detailed overview of the results from [22]. In the following, we use \mathbb{P}_n to denote the underlying probability measure when the change occurs at time n , $n \leq \infty$. We use \mathbb{E}_n to denote the corresponding expectation. We say $p(\alpha) \sim q(\alpha)$ or $p(\alpha) \leq q(\alpha)(1 + o(1))$, as $\alpha \rightarrow 0$, to denote $p(\alpha)/q(\alpha) \rightarrow 1$ and $\lim_{\alpha \rightarrow 0} p(\alpha)/q(\alpha) \leq 1$, respectively, as $\alpha \rightarrow 0$. We use $D(f \parallel g)$ to represent the K-L divergence between the p.d.fs f and g . We assume that the moments of up to third order of all the log likelihood ratios appearing in this paper are finite and positive.

In [22] we considered data-efficient quickest change detection in a single observation sequence. We considered an observation sequence $\{X_n\}$: $\{X_n\}$ are i.i.d. with probability density function (p.d.f.) f_0 before the change point γ , and are i.i.d. with p.d.f. f_1 after the change point γ . A decision maker observes the random variables $\{X_n\}$ over time and has to detect this change in distribution as quickly as possible, subject to constraints on the false alarm rate and the fraction of time observations are taken before change. We now describe the type of policies we consider for this problem.

Let S_n be the indicator random variable such that $S_n = 1$ if X_n is used for decision making, and $S_n = 0$ otherwise. Let

$$\mathcal{I}_n = [S_1, \dots, S_n, X_1^{(S_1)}, \dots, X_n^{(S_n)}],$$

represent the information at time n . Here, $X_i^{(S_i)}$ represents X_i if $S_i = 1$, otherwise X_i is absent from the information vector \mathcal{I}_n . Let τ be a stopping time on the information sequence $\{\mathcal{I}_n\}$, that is, $\mathbb{I}_{\{\tau=n\}}$ is a measurable function of \mathcal{I}_n . Here, \mathbb{I}_F represents the indicator of the event F . For time $n \geq 1$, based on the information vector \mathcal{I}_n , a decision is made whether to *stop and declare change* ($\tau = n$) or to *continue taking observations* ($\tau > n$). If the decision is to *continue*, a decision is made as to whether to *take* or *skip* the observation at time $n + 1$. Thus, S_{n+1} is a function of the information available at time n , i.e.,

$$S_{n+1} = \phi_n(\mathcal{I}_n),$$

where, ϕ_n is the control law at time n . The decision on whether or not to take the first observation is taken without observing $\{X_n\}$. In the absence of a prior information on the distribution of γ , S_1 is typically set to 1, that is the first observation is always taken. A policy for data-efficient QCD is

$$\Psi = \{\tau, \phi_0, \dots, \phi_{\tau-1}\}.$$

To capture the cost of observations used before γ , we proposed a new metric for data-efficiency in minimax settings,

the Pre-change Duty Cycle (PDC). Here we consider its variant that we studied in [23]:

$$\text{PDC}(\Psi) = \limsup_{\gamma \rightarrow \infty} \frac{1}{\gamma} \mathbb{E}_\gamma \left[\sum_{k=1}^{\gamma-1} S_k \right]. \quad (1)$$

We note that $\text{PDC} \leq 1$. If in a policy all the samples are taken, then the PDC for that policy is 1. If every other sample is skipped, then the PDC for that policy is 0.5.

For delay and false alarm we considered the metrics used in [10]: the Worst case Average Detection Delay (WADD)

$$\text{WADD}(\Psi) = \sup_{\gamma \geq 1} \text{ess sup } \mathbb{E}_\gamma [(\tau - \gamma)^+ | \mathcal{I}_{\gamma-1}], \quad (2)$$

and the False Alarm Rate (FAR)

$$\text{FAR}(\Psi) = 1/\mathbb{E}_\infty[\tau]. \quad (3)$$

We considered the following data-efficient minimax formulation.

Problem 2.1:

$$\begin{aligned} & \underset{\Psi}{\text{minimize}} && \text{WADD}(\Psi), \\ & \text{subject to} && \text{FAR}(\Psi) \leq \alpha, \\ & && \text{and} \quad \text{PDC}(\Psi) \leq \beta. \end{aligned} \quad (4)$$

Here $0 \leq \alpha, \beta \leq 1$ are the given constraints.

We also studied the data-efficient minimax formulation where instead of WADD, the following Conditional Average Detection Delay (CADD) metric from [11] is used:

$$\text{CADD}(\Psi) = \sup_{\gamma \geq 1} \mathbb{E}_\gamma [\tau - \gamma | \tau \geq \gamma]. \quad (5)$$

Problem 2.2:

$$\begin{aligned} & \underset{\Psi}{\text{minimize}} && \text{CADD}(\Psi), \\ & \text{subject to} && \text{FAR}(\Psi) \leq \alpha, \\ & && \text{and} \quad \text{PDC}(\Psi) \leq \beta. \end{aligned} \quad (6)$$

Here $0 \leq \alpha, \beta \leq 1$ are given constraints.

We then proposed an algorithm, that we called the DE-CuSum algorithm, and showed that it is asymptotically optimal for both the above problems, for each fixed β , as $\alpha \rightarrow 0$. The DE-CuSum algorithm is defined below.

Algorithm 2.1 (DE-CuSum): Start with $W_0 = 0$ and fix $\mu > 0$, $A > 0$ and $h \geq 0$. For $n \geq 0$ use the following control:

- 1) Take the first observation.
- 2) If an observation is taken then update the statistic using

$$W_{n+1} = (W_n + \log[f_1(X_{n+1})/f_0(X_{n+1})])^{h+},$$

where $(x)^{h+} = \max\{x, -h\}$.

- 3) If $W_n < 0$, skip the next observation and update the statistic using

$$W_{n+1} = \min\{W_n + \mu, 0\}.$$

- 4) Declare change at

$$\tau_w = \inf \{n \geq 1 : W_n > A\}.$$

If $h = 0$, the DE-CuSum statistic W_n never becomes negative and hence reduces to the CuSum statistic [7] and evolves as: $C_0 = 0$, and for $n \geq 0$,

$$C_{n+1} = \max\{0, C_n + \log[f_1(X_{n+1})/f_0(X_{n+1})]\}.$$

The evolution of the DE-CuSum algorithm and the CuSum algorithm for a given sequence of observations is plotted in Fig. 1. If $h = \infty$, the DE-CuSum statistic evolves as follows. Initially the DE-CuSum statistic evolves according to the CuSum statistic till the statistic W_n goes below 0. Once the statistic goes below 0, samples are skipped depending on the undershoot of W_n (this is also the sum of the log likelihood ratio of the observations) and the design parameter μ . Specifically, the statistic is incremented by μ at each time step, and samples are skipped till W_n goes above zero, at which time it is reset to zero. At this point, fresh observations are taken and the process is repeated till the statistic crosses the threshold A , at which time a change is declared. The parameter μ is a substitute for the Bayesian prior ρ that is used in the DE-Shiryayev algorithm described in [21], and is chosen to meet the constraint on the PDC. Thus, the DE-CuSum algorithm is a sequence of SPRTs ([24], [25]) intercepted by “sleep” times controlled by the undershoot and the parameter μ . If $h < \infty$, the number of consecutive samples skipped is bounded by $h/\mu + 1$.




Fig. 1: Typical evolution of the CuSum statistic and the DE-CuSum statistic evaluated using the same set of observations. Note that the CuSum statistic is always greater than the DE-CuSum statistic.

Let τ_C represent the stopping time for the CuSum algorithm. It is well known that the CuSum algorithm is asymptotically optimal for both Problem 2.1 and Problem 2.2, for $\beta = 1$, as $\alpha \rightarrow 0$; see [10], [26] and [27]. However, since all the observations are taken in the CuSum algorithm

$$\text{PDC}(\tau_C) = 1.$$

Thus, the CuSum algorithm is not asymptotically optimal for Problem 2.1 and Problem 2.2, if $\beta < 1$.

We proved the following optimality result for the DE-CuSum algorithm in [22]; also see [23]. Let

$$\tau_- = \inf \left\{ n \geq 1 : \sum_{k=1}^n \log \frac{f_1(X_k)}{f_0(X_k)} < 0 \right\},$$

be the ladder variable. Then W_{τ_-} is the ladder height; see [28] and [25]. Recall that $(x)^{h+} = \max\{x, -h\}$.

Theorem 2.1 ([22], [23]): Let $\mathbb{E}_1[\log[f_1(X_1)/f_0(X_1)]]$ and $\mathbb{E}_\infty[\log[f_0(X_1)/f_1(X_1)]]$ be finite and positive. If $\mu > 0$, $h < \infty$, and $A = |\log \alpha|$, we have

$$\begin{aligned} C_n &\geq W_n \quad \forall n \geq 0, \\ \text{FAR}(\tau_w) &\leq \text{FAR}(\tau_c) \leq \alpha, \\ \text{PDC}(\tau_w) &= \frac{\mathbb{E}_\infty[\tau_-]}{\mathbb{E}_\infty[\tau_-] + \mathbb{E}_\infty[|W_{\tau_-}^{h+}|/\mu]}, \\ \text{CADD}(\tau_w) &\leq \text{CADD}(\tau_c) + K, \\ \text{WADD}(\tau_w) &\leq \text{WADD}(\tau_c) + K. \end{aligned} \quad (7)$$

In the above equation K is a positive constant that is a function of μ , h , and the pre- and post-change distributions, but is not a function of the threshold A . If $h = \infty$, then

$$\text{PDC}(\tau_w) \leq \frac{\mu}{\mu + D(f_0 || f_1)}. \quad (8)$$

Thus, for any fixed threshold, the FAR of the DE-CuSum algorithm is smaller than that of the CuSum algorithm. Thus, we can use the same threshold $A = |\log \alpha|$, as used for the CuSum algorithm, to satisfy the FAR constraint of α . From [10] and [26] it is well known that

$$\begin{aligned} \inf_{\tau: \text{FAR}(\tau) \leq \alpha} \text{CADD}(\tau) &\sim \text{CADD}(\tau_c) \\ &\sim \frac{|\log \alpha|}{D(f_1 || f_0)}, \text{ as } \alpha \rightarrow 0. \end{aligned} \quad (9)$$

Thus, $\frac{|\log \alpha|}{D(f_1 || f_0)}$ is an asymptotic lower bound on the CADD of any stopping time satisfying an FAR constraint of α . Because of the above theorem, the CADD of the DE-CuSum algorithm also achieves this lower bound, for each fixed μ and h (because the delays of the two algorithms are within a constant of each other). Also, since the expression for the PDC is not a function of A , there exists choice of μ and h such that the PDC constraint of β can be satisfied independent of the choice of threshold A . Hence, the DE-CuSum algorithm is asymptotic optimal, for both Problem 2.1 and Problem 2.2, for each fixed β , as $\alpha \rightarrow 0$.

To design the DE-CuSum algorithm to achieve a smaller value of PDC, that is to design the algorithm to drop a larger fraction of samples before change, one has to select a smaller value for the parameter μ . We remark that the assumption that $h < \infty$ is crucial to the proof of the above theorem. In [22] we also showed via simulations that the DE-CuSum algorithm provides a significant gain in performance as compared to the approach of fractional sampling, where the CuSum algorithm is used and the PDC constraint is met by skipping samples randomly, independent of the observation process.

III. PROBLEM FORMULATION FOR OUTLYING SEQUENCE DETECTION IN SENSOR NETWORKS

We have a sensor network consisting of L sensors and a central decision maker called the fusion center. The sensors are indexed by the index $\ell \in \{1, \dots, L\}$. At sensor ℓ the sequence $\{X_{n,\ell}\}_{n \geq 1}$ is observed, where n is the time index. At γ , the distribution of $\{X_{n,\ell}\}$ in a subset $\kappa = \{k_1, k_2, \dots, k_m\} \subset$

$\{1, 2, \dots, L\}$ of the sensor nodes changes, from $f_{0,\ell}$ to say $f_{1,\ell}$. The random variables $\{X_{n,\ell}\}$ are independent across indices n and ℓ conditioned on γ and the affected subset κ . The distributions $f_{0,\ell}$ and $f_{1,\ell}$ are assumed to be known, but neither the affected subset κ nor its size m is known.

Let $S_{n,\ell}$ be the indicator random variable such that

$$S_{n,\ell} = \begin{cases} 1 & \text{if } X_{n,\ell} \text{ is used for decision making at sensor } \ell \\ 0 & \text{otherwise.} \end{cases}$$

Let $\phi_{n,\ell}$ be the observation control law at sensor ℓ , i.e.,

$$S_{n+1,\ell} = \phi_{n,\ell}(\mathcal{I}_{n,\ell}),$$

where $\mathcal{I}_{n,\ell} = [S_{1,\ell}, \dots, S_{n,\ell}, X_{1,\ell}^{(S_{1,\ell})}, \dots, X_{n,\ell}^{(S_{n,\ell})}]$. Here, $X_{n,\ell}^{(S_{n,\ell})} = X_{1,\ell}$ if $S_{1,\ell} = 1$, otherwise $X_{1,\ell}$ is absent from the information vector $\mathcal{I}_{n,\ell}$. Thus, the decision to take or skip a sample at sensor ℓ is based on its past information. Let

$$\mathcal{I}_n = \{\mathcal{I}_{n,1}, \dots, \mathcal{I}_{n,L}\}$$

be the information available at time n across the sensor network.

Also let

$$Y_{n,\ell} = g_{n,\ell}(\mathcal{I}_{n,\ell})$$

be the information transmitted from sensor ℓ to the fusion center. If no information is transmitted to the fusion center, then $Y_{n,\ell} = \text{NULL}$, which is treated as zero at the fusion center. Here, $g_{n,\ell}$ is the transmission control law at sensor ℓ . Let

$$\mathbf{Y}_n = \{Y_{n,1}, \dots, Y_{n,L}\}$$

be the information received at the fusion center at time n , and let τ be a stopping time on the sequence $\{\mathbf{Y}_n\}$.

Let

$$\phi_n = \{\phi_{n,1}, \dots, \phi_{n,L}\}$$

denote the observation control law at time n , and let

$$\mathbf{g}_n = \{g_{n,1}, \dots, g_{n,L}\}$$

denote the transmission control law at time n . For data-efficient QCD in sensor networks we consider the policy of type II defined as

$$\Pi = \{\tau, \{\phi_0, \dots, \phi_{\tau-1}\}, \{\mathbf{g}_1, \dots, \mathbf{g}_\tau\}\}.$$

The PDC_ℓ , the PDC for sensor ℓ , is defined as

$$\text{PDC}_\ell(\Pi) = \limsup_{\gamma \rightarrow \infty} \frac{1}{\gamma} \mathbb{E}_\infty \left[\sum_{k=1}^{\gamma-1} S_{k,\ell} \right]. \quad (10)$$

Thus, PDC_ℓ is the fraction of time observations are taken before change at sensor ℓ .

To capture the cost of communication between each sensor and the fusion center before change, we propose the Pre-Change Transmission Cost (PTC) metric. We define

$$T_{n,\ell} = \begin{cases} 1 & \text{if } Y_{n,\ell} \neq \text{NULL, i.e., some information} \\ & \text{is transmitted to the fusion center} \\ 0 & \text{otherwise.} \end{cases}$$

The Pre-change Transmission Cost at sensor ℓ (PTC_ℓ) is defined as

$$\text{PTC}_\ell(\Pi) = \limsup_{\gamma \rightarrow \infty} \frac{1}{\gamma} \mathbb{E}_\infty \left[\sum_{k=1}^{\gamma-1} T_{k,\ell} \right]. \quad (11)$$

If in a policy every sample is taken and some information is transmitted at every time slot at all the sensors, then for that policy $\text{PDC}_\ell = \text{PTC}_\ell = 1, \forall \ell$. If transmissions happen from the sensors only in every alternate time slots, then $\text{PTC}_\ell = 0.5, \forall \ell$.

The objective here is to solve the following extensions of Problem 2.1 and Problem 2.2:

Problem 3.1:

$$\begin{aligned} & \underset{\Pi}{\text{minimize}} && \text{WADD}(\Pi), \\ & \text{subject to} && \text{FAR}(\Pi) \leq \alpha, \\ & && \text{PDC}_\ell(\Pi) \leq \beta_\ell, \text{ for } \ell = 1, \dots, L, \\ & \text{and} && \text{PTC}_\ell(\Pi) \leq \sigma_\ell, \text{ for } \ell = 1, \dots, L, \end{aligned} \quad (12)$$

where $0 \leq \alpha, \beta_\ell, \sigma_\ell \leq 1$, for $\ell = 1, \dots, L$, are given constraints, and

Problem 3.2:

$$\begin{aligned} & \underset{\Pi}{\text{minimize}} && \text{CADD}(\Pi), \\ & \text{subject to} && \text{FAR}(\Pi) \leq \alpha, \\ & && \text{PDC}_\ell(\Pi) \leq \beta_\ell, \text{ for } \ell = 1, \dots, L, \\ & \text{and} && \text{PTC}_\ell(\Pi) \leq \sigma_\ell, \text{ for } \ell = 1, \dots, L, \end{aligned} \quad (13)$$

where $0 \leq \alpha, \beta_\ell, \sigma_\ell \leq 1$, for $\ell = 1, \dots, L$, are given constraints.

The asymptotic lower bound developed in [26] can be specialized to the sensor network setting considered. Let

$$\Delta_\alpha = \{\Pi : \text{FAR}(\Pi) \leq \alpha\}.$$

Theorem 3.1 ([26]): If the outlying subset post-change is $\kappa = \{k_1, k_2, \dots, k_m\}$, then as $\alpha \rightarrow 0$,

$$\inf_{\Pi \in \Delta_\alpha} \text{CADD}(\Pi) \geq \frac{|\log \alpha|}{\sum_{i=1}^m D(f_{1,k_i} || f_{0,k_i})} (1 + o(1)). \quad (14)$$

Since $\text{WADD}(\Pi) \geq \text{CADD}(\Pi)$, we also have as $\alpha \rightarrow 0$,

$$\inf_{\Pi \in \Delta_\alpha} \text{WADD}(\Pi) \geq \frac{|\log \alpha|}{\sum_{i=1}^m D(f_{1,k_i} || f_{0,k_i})} (1 + o(1)). \quad (15)$$

IV. DATA-EFFICIENT ALGORITHMS FOR OUTLYING SEQUENCE DETECTION IN SENSOR NETWORKS

In this section we propose two algorithms that can be used to detect the outlying sequences in a data-efficient way in a sensor network. In both the algorithms the DE-CuSum algorithm (Algorithm 2.1) is used locally at each sensor. In the rest of the paper we use $W_{n,\ell}$ to denote the DE-CuSum statistic at sensor ℓ .

A. The DE-Censor-Max Algorithm

In the DE-Censor-Max algorithm, the DE-CuSum algorithm is used at each sensor ℓ . If the DE-CuSum statistic $W_{n,\ell}$ at a sensor is above a threshold D_ℓ , then the statistic is transmitted to the fusion center. A change is declared at the fusion center, if the maximum of the transmitted statistics from all the sensors is larger than another threshold A . Mathematically, the DE-Censor-Max algorithm is described as follows.

Algorithm 4.1 (DE-Censor-Max: Π_{DCM}): Start with $W_{0,\ell} = 0, \forall \ell$. Fix $\mu_\ell > 0, h_\ell \geq 0, D_\ell \geq 0$ and $A \geq 0$. For $n \geq 0$ use the following control:

- 1) Use the DE-CuSum algorithm at each sensor ℓ , i.e., update the statistics $\{W_{n,\ell}\}_{\ell=1}^L$ for $n \geq 1$ using

$$S_{n+1,\ell} = 1 \text{ only if } W_{n,\ell} \geq 0$$

$$W_{n+1,\ell} = \begin{cases} \min\{W_{n,\ell} + \mu_\ell, 0\} & \text{if } S_{n+1,\ell} = 0 \\ \left(W_{n,\ell} + \log \frac{f_{1,\ell}(X_{n+1,\ell})}{f_{0,\ell}(X_{n+1,\ell})}\right)^{h+} & \text{if } S_{n+1,\ell} = 1, \end{cases}$$
 where $(x)^{h+} = \max\{x, -h\}$.
- 2) Transmit

$$Y_{n,\ell} = W_{n,\ell} \mathbb{I}_{\{W_{n,\ell} > D_\ell\}}.$$

- 3) At the fusion center stop at

$$\tau_{\text{DCM}} = \inf\{n \geq 1 : \max_{\ell \in \{1, \dots, L\}} Y_{n,\ell} > A\}.$$

With $D_\ell = 0$ and $h_\ell = 0, \forall \ell$, the DE-CuSum algorithm at each sensor reduces to the CuSum algorithm, and $Y_{n,\ell} = W_{n,\ell} \forall n, \ell$. In this case, the DE-Censor-Max algorithm reduces to the MAX algorithm proposed in [16].

We will show in the Section V that when exactly one of the L sensor is affected post-change, then this algorithm is uniformly asymptotically optimal for both Problem 3.1 and Problem 3.2 (achieves the lower bound provided in Theorem 3.1 for each κ), for each fixed $\{\beta_\ell\}$ and $\{\sigma_\ell\}$, as $\alpha \rightarrow 0$.

B. The DE-Censor-Sum Algorithm

Although the DE-Censor-Max algorithm is asymptotically optimal, we will show in Section VI that it performs poorly when the size of the outlying subset is large. To address this deficiency, we propose the DE-Censor-Sum algorithm.

In the DE-Censor-Sum algorithm, the DE-CuSum algorithm is used at each sensor. If the DE-CuSum statistic at a sensor is above a threshold, then the statistic is transmitted to the fusion center. A change is declared at the fusion center, if the sum of the transmitted statistics from all the sensors is larger than another threshold. Mathematically, the DE-Censor-Sum algorithm is described as follows.

Algorithm 4.2 (DE-Censor-Sum: Π_{DCS}): Start with $W_{0,\ell} = 0, \forall \ell$. Fix $\mu_\ell > 0, h_\ell \geq 0, D_\ell \geq 0$ and $A \geq 0$. For $n \geq 0$ use the following control:

- 1) Use the DE-CuSum algorithm at each sensor ℓ , i.e., update the statistics $\{W_{n,\ell}\}_{\ell=1}^L$ for $n \geq 1$ using

$$S_{n+1,\ell} = 1 \text{ only if } W_{n,\ell} \geq 0$$

$$W_{n+1,\ell} = \begin{cases} \min\{W_{n,\ell} + \mu_\ell, 0\} & \text{if } S_{n+1,\ell} = 0 \\ \left(W_{n,\ell} + \log \frac{f_{1,\ell}(X_{n+1,\ell})}{f_{0,\ell}(X_{n+1,\ell})}\right)^{h+} & \text{if } S_{n+1,\ell} = 1, \end{cases}$$

where $(x)^{h+} = \max\{x, -h\}$.

2) Transmit

$$Y_{n,\ell} = W_{n,\ell} \mathbb{I}_{\{W_{n,\ell} > D_\ell\}}.$$

3) At the fusion center stop at

$$\tau_{\text{DCS}} = \inf\{n \geq 1 : \sum_{\ell \in \{1, \dots, L\}} Y_{n,\ell} > A\}.$$

With $D_\ell = 0$ and $h_\ell = 0$, $\forall \ell$, the DE-CuSum algorithm at each sensor reduces to the CuSum algorithm, and $Y_{n,\ell} = W_{n,\ell}$ $\forall n, \ell$. In this case, the DE-Censor-Sum algorithm reduces to the N_{sum} algorithm proposed in [17]. If $h_\ell = 0 \forall \ell$ and $D > 0$, the DE-Censor-Sum algorithm reduces to the N_{hard} algorithm proposed in [18]. The DE-Censor-Sum algorithm can easily be modified to obtain data-efficient extensions of other algorithms proposed in [18].

We will provide a detailed performance analysis of the DE-Censor-Sum algorithm using which the threshold A and the parameter h_ℓ and μ_ℓ can be selected. We will use the performance analysis to show that, under the additional assumption a result in [17], the DE-Censor-Sum algorithm is also uniformly asymptotically optimal for both Problem 3.1 and Problem 3.2 (achieves the lower bound provided in Theorem 3.1 for each κ), for each fixed $\{\beta_\ell\}$ and $\{\sigma_\ell\}$, as $\alpha \rightarrow 0$.

V. ASYMPTOTIC OPTIMALITY OF THE DE-CENSOR-MAX ALGORITHM

In this section we now show the asymptotic optimality of the DE-Censor-Max algorithm.

We define the ladder variable [28] corresponding to sensor ℓ :

$$\tau_{\ell-} = \inf \left\{ n \geq 1 : \sum_{k=1}^n \log \frac{f_{1,\ell}(X_{k,\ell})}{f_{0,\ell}(X_{k,\ell})} < 0 \right\},$$

and note that $W_{\tau_{\ell-}}$ is the ladder height. Also, let

$$U_{D_\ell} = \left\{ \# \text{ times} : \sum_{k=1}^n \log \frac{f_{1,\ell}(X_{k,\ell})}{f_{0,\ell}(X_{k,\ell})} > D_\ell \text{ before it's } < 0 \right\}.$$

Thus, U_{D_ℓ} is the number of times the random walk $\sum_{k=1}^n \log \frac{f_{1,\ell}(X_{k,\ell})}{f_{0,\ell}(X_{k,\ell})}$ is above D_ℓ before hitting 0. We note that $U_{D_\ell} = 0$ with a positive probability. We note that U_{D_ℓ} is also the number of times the DE-CuSum statistic $W_{n,\ell}$ is above D_ℓ before hitting 0.

Also for x_ℓ real define

$$\tau_{w,\ell}(x_\ell, A) = \inf\{n \geq 1 : W_{n,\ell} > A; \text{ with } W_{0,\ell} = x_\ell\}. \quad (16)$$

Thus, $\tau_{w,\ell}(x_\ell, A)$ is the time for the DE-CuSum statistic at sensor ℓ to cross the threshold A starting with the initial value x_ℓ . It follows from Lemma 5 in [22] that

$$\mathbb{E}_1[\tau_{w,\ell}(x_\ell, A)] \leq \mathbb{E}_1[\tau_{w,\ell}(0, A)] + \lceil h_\ell / \mu_\ell \rceil, \quad \forall \ell. \quad (17)$$

Theorem 5.1: Let

$$0 < D(f_{1,\ell} \parallel f_{0,\ell}) < \infty \text{ and } 0 < D(f_{0,\ell} \parallel f_{1,\ell}) < \infty \quad \forall \ell.$$

Let $\mu_\ell > 0$, $h_\ell < \infty$, $\forall \ell$, $D_\ell \geq 0$, and $A = \log \frac{L}{\alpha}$. If the change occurs in the stream ℓ^* , then we have

$$\text{FAR}(\Pi_{\text{DCM}}) \leq \alpha,$$

$$\text{PDC}_\ell(\Pi_{\text{DCM}}) = \frac{\mathbb{E}_\infty[\tau_{\ell-}]}{\mathbb{E}_\infty[\tau_{\ell-}] + \mathbb{E}_\infty[\lceil |W_{\tau_{\ell-}}^{h_\ell+}| / \mu_\ell \rceil]}, \quad \forall \ell,$$

$$\text{PTC}_\ell(\Pi_{\text{DCM}}) = \frac{\mathbb{E}_\infty[U_{D_\ell}]}{\mathbb{E}_\infty[\tau_{\ell-}] + \mathbb{E}_\infty[\lceil |W_{\tau_{\ell-}}^{h_\ell+}| / \mu_\ell \rceil]}, \quad \forall \ell,$$

$$\text{WADD}(\Pi_{\text{DCM}}) \leq \frac{|\log \alpha|}{D(f_{1,\ell^*} \parallel f_{0,\ell^*})} (1 + o(1)) \text{ as } \alpha \rightarrow 0. \quad (18)$$

If $h_\ell = \infty$, $\forall \ell$, then

$$\text{PDC}_\ell(\Pi_{\text{DCM}}) \leq \frac{\mu_\ell}{\mu_\ell + D(f_{0,\ell} \parallel f_{1,\ell})} \quad \forall \ell. \quad (19)$$

Proof: The FAR result follows from Theorem 1 of [16] because the DE-CuSum statistic at each sensor is always greater than the CuSum statistic computed using the same set of observations; see Theorem 2.1.

The results on PDC_ℓ follows from Theorem 2.1. The idea behind the proof is to define an on-off renewal process, with on times distributed according to the time for which the DE-CuSum statistic $W_{n,\ell}$ at sensor ℓ is above 0, and the off times distributed according to the sojourn of the DE-CuSum statistic $W_{n,\ell}$ at sensor ℓ is above 0. The result then follows from the renewal reward theorem. The results on PTC_ℓ follows also from the renewal reward theorem and the arguments are almost identical to those provide for PDC_ℓ .

The delay result holds because after change the max of statistics is greater than the statistics in which the change has taken place. Thus, the delay of the DE-Censor-Max algorithm is bounded from above by the delay of the DE-CuSum algorithm when applied to the outlying sensor. Mathematically, the argument is as follows.

We obtain an upper bound on $\mathbb{E}_\gamma[(\tau_{\text{DCM}} - \gamma)^+ | \mathcal{I}_{\gamma-1}]$ that is not a function of γ and the conditioning $\mathcal{I}_{\gamma-1}$, and that scales as the lower bound in Theorem 3.1. The theorem is then established if we then take the essential supremum and then the supremum over γ .

Let $\mathcal{I}_{\gamma-1} = \mathbf{i}_{\gamma-1}$ be such that $W_{\gamma-1,\ell} = x_\ell$, $x_\ell \in [-h_\ell, \infty)$. We first note that for $A > \max_\ell D_\ell$,

$$\mathbb{E}_\gamma[(\tau_{\text{DCM}} - \gamma)^+ | \mathcal{I}_{\gamma-1} = \mathbf{i}_{\gamma-1}] \leq \mathbb{E}_1[\tau_{w,\ell^*}(x_{\ell^*}, A)], \quad (20)$$

where $\tau_{w,\ell}(x_\ell, A)$ is as defined in (16). Then from (17) we have

$$\begin{aligned} \mathbb{E}_\gamma[(\tau_{\text{DCM}} - \gamma)^+ | \mathcal{I}_{\gamma-1} = \mathbf{i}_{\gamma-1}] \\ \leq \mathbb{E}_1[\tau_{w,\ell^*}(x_{\ell^*}, A)] \leq \mathbb{E}_1[\tau_{w,\ell^*}(0, A)] + \lceil h_{\ell^*} / \mu_{\ell^*} \rceil. \end{aligned} \quad (21)$$

The result now follows from the proof of Theorem 2.1 on the DE-CuSum algorithm because the change is assumed to affect the sensor ℓ^* and $h_{\ell^*} < \infty$. ■

Since $\text{CADD} \leq \text{WADD}$, we also have under the same assumptions as in Theorem 5.1

$$\text{CADD}(\Pi_{\text{DCM}}) \leq \frac{|\log \alpha|}{D(f_{1,\ell^*} \parallel f_{0,\ell^*})} (1 + o(1)), \text{ as } \alpha \rightarrow 0. \quad (22)$$

From Theorem 3.1, the WADD, and hence the CADD performance of the DE-Censor-Max algorithm is the best one can do when the change affects the stream ℓ^* , for given $\{\beta_\ell\}$ and $\{\sigma_\ell\}$, as $\alpha \rightarrow 0$. Also, the PDC_ℓ and the PTC_ℓ performances do not depend on the threshold A , thus the constraints $\{\beta_\ell\}$ and $\{\sigma_\ell\}$ can be satisfied independent of the FAR constraint α . Hence, the DE-Censor-Max algorithm is asymptotically optimal when the change affects exactly one stream, for both Problem 3.1 and Problem 3.2, for each given $\{\beta_\ell\}$ and $\{\sigma_\ell\}$, as $\alpha \rightarrow 0$.

A. Performance Analysis of the DE-Censor-Sum Algorithm

In this section we provide the performance analysis of the DE-Censor-Sum algorithm and comment on its asymptotic optimality. In Theorem 5.2 to be proved below, the delay analysis of the DE-Censor-Sum depends on the delay analysis of the DE-All algorithm that we studied in [23]. We first define the DE-All algorithm.

In the DE-All algorithm, the DE-CuSum algorithm (see Algorithm 2.1) is used at each sensor, and a “1” is transmitted each time the DE-CuSum statistic at any sensor is above a threshold. A change is declared the first time a “1” is received at the fusion center from *all* the sensors at the same time.

Let

$$d_\ell = \frac{D(f_{1,\ell} \parallel f_{0,\ell})}{\sum_{k=1}^L D(f_{1,k} \parallel f_{0,k})}.$$

Algorithm 5.1 (DE – All: Π_{All}): Start with $W_{0,\ell} = 0 \forall \ell$. Fix $\mu_\ell > 0$, $h_\ell \geq 0$, and $A \geq 0$. For $n \geq 0$ use the following control:

- 1) Use the DE-CuSum algorithm at each sensor ℓ , i.e., update the statistics $\{W_{n,\ell}\}_{\ell=1}^L$ for $n \geq 1$ using

$$S_{n+1,\ell} = 1 \text{ only if } W_{n,\ell} \geq 0$$

$$W_{n+1,\ell} = \begin{cases} \min\{W_{n,\ell} + \mu_\ell, 0\} & \text{if } S_{n+1,\ell} = 0 \\ \left(W_{n,\ell} + \log \frac{f_{1,\ell}(X_{n+1,\ell})}{f_{0,\ell}(X_{n+1,\ell})}\right)^{h+} & \text{if } S_{n+1,\ell} = 1, \end{cases}$$

where $(x)^{h+} = \max\{x, -h\}$.

- 2) Transmit

$$Y_{n,\ell} = \mathbb{I}_{\{W_{n,\ell} > d_\ell A\}}.$$

- 3) At the fusion center stop at

$$\tau_{\text{DE-All}} = \inf\{n \geq 1 : Y_{n,\ell} = 1 \text{ for all } \ell \in \{1, \dots, L\}\}.$$

We now provide the main result of this section. For that we define another variable:

$$\tau_{C,\ell}(x_\ell, A) = \inf\{n \geq 1 : C_{n,\ell} > A; \text{ with } C_{0,\ell} = x_\ell\}. \quad (23)$$

Thus, $\tau_{C,\ell}(x_\ell, A)$ is the time for the CuSum statistic at sensor ℓ $C_{n,\ell}$, to cross the threshold A starting with the initial value x_ℓ .

Theorem 5.2: Let

$$0 < D(f_{1,\ell} \parallel f_{0,\ell}) < \infty \text{ and } 0 < D(f_{0,\ell} \parallel f_{1,\ell}) < \infty \quad \forall \ell.$$

Let $\mu_\ell > 0$, $h_\ell < \infty$, $\forall \ell$, $D_\ell \geq 0$, and $A = L \log \frac{L}{\alpha}$. If the change affects the subset κ of streams, then we have

$$\begin{aligned} \text{FAR}(\Pi_{\text{DCS}}) &\leq \alpha, \\ \text{PDC}_\ell(\Pi_{\text{DCS}}) &= \frac{\mathbb{E}_\infty[\tau_{\ell-}]}{\mathbb{E}_\infty[\tau_{\ell-}] + \mathbb{E}_\infty[|W_{\tau_{\ell-}}^{h_\ell+}|/\mu_\ell]}, \quad \forall \ell \\ \text{PTC}_\ell(\Pi_{\text{DCS}}) &= \frac{\mathbb{E}_\infty[U_{D_\ell}]}{\mathbb{E}_\infty[\tau_{\ell-}] + \mathbb{E}_\infty[|W_{\tau_{\ell-}}^{h_\ell+}|/\mu_\ell]}, \quad \forall \ell \\ \text{WADD}(\Pi_{\text{DCS}}) &\leq \frac{A}{\sum_{i=1}^m D(f_{1,k_i} \parallel f_{0,k_i})} (1 + o(1)) \text{ as } A \rightarrow \infty. \end{aligned} \quad (24)$$

If $h_\ell = \infty$, $\forall \ell$, then

$$\text{PDC}_\ell(\Pi_{\text{DCS}}) \leq \frac{\mu_\ell}{\mu_\ell + D(f_{0,\ell} \parallel f_{1,\ell})} \quad \forall \ell. \quad (25)$$

Proof: The proofs on PDC_ℓ and PTC_ℓ are identical to that provided in the Theorem 5.1.

For the FAR note that

$$\left\{ \sum_{\ell=1}^L W_{n,\ell} > A \right\} \subset \left\{ \max_{\ell \in \{1, \dots, L\}} W_{n,\ell} > \frac{A}{L} \right\}.$$

For simplicity we write $\Pi_{\text{DCS}}(A)$ to represent DE-Censor-Sum algorithm when the threshold used at the fusion center is A . Similarly we use $\Pi_{\text{DCM}}(A/L)$ to represent DE-Censor-Max algorithm when the threshold used at the fusion center is A/L . Then the above subset relation implies

$$\text{FAR}(\Pi_{\text{DCS}}(A)) \leq \text{FAR}(\Pi_{\text{DCM}}(A/L)).$$

The FAR result follows because from Theorem 5.1 we have that

$$\text{FAR}(\Pi_{\text{DCM}}(A/L)) \leq \alpha \quad \text{if} \quad A/L = \log L/\alpha.$$

For the WADD analysis, let $\tau_{\text{DCS}}(\kappa)$ denote the DE-Censor-Sum algorithm applied to only the streams in the affected subset κ . Further let $\mathcal{I}_{\gamma-1}(\kappa)$ denote the information in the affected streams. Then

$$\begin{aligned} \mathbb{E}_\gamma[(\tau_{\text{DCS}} - \gamma)^+ | \mathcal{I}_{\gamma-1} = \mathbf{i}_{\gamma-1}] &\leq \mathbb{E}_\gamma[(\tau_{\text{DCS}}(\kappa) - \gamma)^+ | \mathcal{I}_{\gamma-1} = \mathbf{i}_{\gamma-1}] \\ &= \mathbb{E}_\gamma[(\tau_{\text{DCS}}(\kappa) - \gamma)^+ | \mathcal{I}_{\gamma-1}(\kappa) = \mathbf{i}_{\gamma-1}(\kappa)]. \end{aligned} \quad (26)$$

In the above equation the last equality is true because the observations across the streams are independent conditioned on the change point. Because of the above inequality, we can assume that the change affects all the subsets at the same time, i.e., $\kappa = \{1, \dots, L\}$.

Now note that any fixed A (see Algorithm 5.1)

$$\{W_{n,\ell} > d_\ell A, \forall \ell\} \subset \left\{ \sum_{\ell} W_{n,\ell} > \sum_{\ell} d_\ell A = A \right\}.$$

Hence, for A sufficiently large and from the proof of Theorem 6.1 in [23], we have

$$\begin{aligned} \mathbb{E}_\gamma[(\tau_{\text{DCS}} - \gamma)^+ | \mathcal{I}_{\gamma-1} = \mathbf{i}_{\gamma-1}] &\leq \mathbb{E}_\gamma[(\tau_{\text{DE-All}} - \gamma)^+ | \mathcal{I}_{\gamma-1} = \mathbf{i}_{\gamma-1}] \\ &\leq \mathbb{E}_1 \left[\max_{1 \leq \ell \leq L} \tau_{C,\ell} \right] + \text{constant}. \end{aligned} \quad (27)$$

The proof of the theorem is now complete because we can now take ess sup and then sup over γ on the left-hand side. Then, from [14] it follows that $\mathbb{E}_1[\max_{1 \leq \ell \leq L} \tau_{C,\ell}]$ grows in the order $\frac{A}{\sum_{\ell=1}^L D(f_{1,\ell} || f_{0,\ell})}$. This when applied to the affected subset κ gives us the desired result on the WADD from (26). ■

Note that the above theorem does not imply the asymptotic optimality of the DE-Censor-Sum algorithm, mainly due to the fact that the choice of the threshold is conservative. It only gives a delay bound of L times that of the lower bound in Theorem 3.1. However, if the threshold can be set to be of the order $\log 1/\alpha$ to satisfy the FAR constraint, then the above theorem establishes the uniform asymptotic optimality of the DE-Censor-Sum algorithm for each possible post-change distribution. It is shown in [17] that such a result is indeed true, and therefore under the conditions of Theorem 5.2 above, the DE-Censor-Sum algorithm is uniformly asymptotically optimal, for each possible κ , for each fixed $\{\beta_\ell\}$ and $\{\sigma_\ell\}$, as $\alpha \rightarrow 0$.

VI. NUMERICAL RESULTS

We first compare the performance of the DE-Censor-Sum algorithm, the DE-Censor-Max algorithm and the Oracle CuSum algorithm as a function of the number of affected stream. The Oracle CuSum algorithm uses all of the sensor observations, and it knows the indices of the outlying sequences. We plot the CADD versus the number of affected stream comparison in Fig. 2 for the parameters: $\text{FAR} = 10^{-3}$, $L = 100$, $f_{0,\ell} = f_0 = \mathcal{N}(0, 1)$, $\forall \ell$, $f_{1,\ell} = f_1 = \mathcal{N}(0.5, 1)$, $\forall \ell$, and for the $\{\text{PDC}_\ell\}$ and $\{\text{PTC}_\ell\}$ constraints of $\beta_\ell = \sigma_\ell = 0.5$, $\forall \ell$. We also set the local thresholds $D_\ell = 0$, $\forall \ell$. In the figure




Fig. 2: Comparison of DE-Censor-Sum algorithm, the DE-Censor-Max algorithm and the Oracle CuSum algorithm as a function of the number of outlying streams.

we see that the DE-Censor-Max scheme outperforms the DE-Censor-Sum scheme when the number of affected streams is small. This is because the former is optimal when the number of affected stream is exactly one. However, when the number of affected streams is large, the DE-Censor-Sum algorithm outperforms the DE-Censor-Max algorithm. We note that this is consistent with the observations made in [17] regarding the comparison between the MAX and SUM algorithms.

In Fig. 3 we compare the CADD vs FAR performance of the DE-Censor-Sum algorithm with the fractional sampling scheme for $L = 10$, $f_{0,\ell} = f_0 = \mathcal{N}(0, 1)$, $\forall \ell$, $f_{1,\ell} = f_1 = \mathcal{N}(0.2, 1)$, $\forall \ell$, and for the $\{\text{PDC}_\ell\}$ and $\{\text{PTC}_\ell\}$ constraints of $\beta_\ell = \sigma_\ell = 0.5$, $\forall \ell$. We consider the post-change scenario when $m = 7$. We restrict our numerical study to the comparison of the CADD performance. Similar comparison can be obtained for the WADD as well.

In the fractional sampling scheme, the CuSum algorithm is used at each sensor, and samples are skipped based on the outcome of a sequence of fair coin tosses, independent of the observation process. If an observation is taken at a sensor, the CuSum statistic is transmitted to the fusion center. Thus, achieving the constraints on the $\{\text{PDC}_\ell\}$ and $\{\text{PTC}_\ell\}$, $\forall \ell$. At the fusion center a change is declared the first time the sum of the CuSum statistics from all the sensors crosses a threshold. At the fusion center, in the absence of any transmission from a sensor, its CuSum statistics from the last time instant is used to compute the sum. For the DE-Censor-Sum algorithm, we set $D_\ell = 0$, $\{h_\ell = h = 10\}$, $\forall \ell$, and use the approximation (19) to select μ_ℓ . This ensures that the $\{\text{PDC}_\ell\}$ and $\{\text{PTC}_\ell\}$ constraints are satisfied for the DE-Censor-Sum algorithm. In the figure we see that the DE-Censor-Sum algorithm provides a significant gain in performance as compared to the approach of fractional sampling.




Fig. 3: Comparison of the DE-Censor-Sum algorithm with the fractional sampling scheme.

VII. CONCLUSIONS

In this paper we proposed two data-efficient algorithms, the DE-Censor-Max algorithm and the DE-Censor-Sum algorithm, for quickest outlying sequence detection in sensor networks when the subset of affected sensors post-change is unknown to the decision maker. We provided a detailed performance analysis of these algorithms and compared their performance as a function of the number of affected sensors. We showed that the DE-Censor-Max algorithm is asymptotically optimal for the proposed formulations if exactly one sensor is affected post-change. Also, if the threshold can be appropriately selected, then the DE-Censor-Sum algorithm is asymptotically optimal, for every possible post-change scenario. We also showed via simulations that the DE-Censor-Max algorithm performs better if the number of affected sensors is small. Moreover, our

algorithms for observation control provide significant benefit over the approach of fractional sampling.

- [28] M. Woodroffe, *Nonlinear Renewal Theory in Sequential Analysis*. CBMS-NSF regional conference series in applied mathematics, SIAM, 1982.

REFERENCES

- [1] H. V. Poor and O. Hadjiladis, *Quickest Detection*. Cambridge University Press, 2009.
- [2] A. G. Tartakovsky, I. V. Nikiforov, and M. Basseville, *Sequential Analysis: Hypothesis Testing and Change-Point Detection*. Statistics, CRC Press, 2014.
- [3] M. Basseville and I. V. Nikiforov, *Detection of Abrupt Changes: Theory and Application*. Englewood Cliffs, NJ: Prentice Hall, 1993.
- [4] V. V. Veeravalli and T. Banerjee, *Quickest Change Detection*. Elsevier: E-reference Signal Processing, 2013. <http://arxiv.org/abs/1210.5552>.
- [5] W. A. Shewhart, *Economic control of quality of manufactured product*. American Society for Quality Control, 1931.
- [6] M. A. Girshick and H. Rubin, "A bayes approach to a quality control model," *Ann. Math. Statist.*, vol. 23, no. 1, pp. 114–125, 1952.
- [7] E. S. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, pp. 100–115, June 1954.
- [8] A. N. Shiryaev, "On optimum methods in quickest detection problems," *Theory of Prob and App.*, vol. 8, pp. 22–46, 1963.
- [9] A. N. Shiryaev, *Optimal Stopping Rules*. New York: Springer-Verlag, 1978.
- [10] G. Lorden, "Procedures for reacting to a change in distribution," *Ann. Math. Statist.*, vol. 42, pp. 1897–1908, Dec. 1971.
- [11] M. Pollak, "Optimal detection of a change in distribution," *Ann. Statist.*, vol. 13, pp. 206–227, Mar. 1985.
- [12] G. V. Moustakides, "Optimal stopping times for detecting changes in distributions," *Ann. Statist.*, vol. 14, pp. 1379–1387, Dec. 1986.
- [13] V. V. Veeravalli, "Decentralized quickest change detection," *IEEE Trans. Inf. Theory*, vol. 47, pp. 1657–1665, May 2001.
- [14] Y. Mei, "Information bounds and quickest change detection in decentralized decision systems," *IEEE Trans. Inf. Theory*, vol. 51, pp. 2669–2681, July 2005.
- [15] A. G. Tartakovsky and V. V. Veeravalli, "Asymptotically optimal quickest change detection in distributed sensor systems," *Sequential Analysis*, vol. 27, pp. 441–475, Oct. 2008.
- [16] A. G. Tartakovsky and V. V. Veeravalli, "An efficient sequential procedure for detecting changes in multichannel and distributed systems," in *IEEE International Conference on Information Fusion*, vol. 1, (Annapolis, MD), pp. 41–48, July 2002.
- [17] Y. Mei, "Efficient scalable schemes for monitoring a large number of data streams," *Biometrika*, vol. 97, pp. 419–433, Apr. 2010.
- [18] Y. Mei, "Quickest detection in censoring sensor networks," in *IEEE International Symposium on Information Theory (ISIT)*, pp. 2148–2152, Aug. 2011.
- [19] L. Zacharias and R. Sundaresan, "Decentralized sequential change detection using physical layer fusion," *IEEE Trans. Wireless Commun.*, vol. 7, pp. 4999–5008, Dec. 2008.
- [20] T. Banerjee, V. Sharma, V. Kavitha, and A. K. Jayaprakasam, "Generalized analysis of a distributed energy efficient algorithm for change detection," *IEEE Trans. Wireless Commun.*, vol. 10, pp. 91–101, Jan. 2011.
- [21] T. Banerjee and V. V. Veeravalli, "Data-efficient quickest change detection with on-off observation control," *Sequential Analysis*, vol. 31, pp. 40–77, Feb. 2012.
- [22] T. Banerjee and V. V. Veeravalli, "Data-efficient quickest change detection in minimax settings," *IEEE Trans. Inf. Theory*, vol. 59, pp. 6917 – 6931, Oct. 2013.
- [23] T. Banerjee and V. V. Veeravalli, "Data-efficient minimax quickest change detection in a decentralized system," *Submitted to Sequential Analysis*, 2014. <http://arxiv.org/abs/1408.4747>.
- [24] A. Wald and J. Wolfowitz, "Optimum character of the sequential probability ratio test," *Ann. Math. Statist.*, vol. 19, no. 3, pp. 326–339, 1948.
- [25] D. Siegmund, *Sequential Analysis: Tests and Confidence Intervals*. Springer series in statistics, Springer-Verlag, 1985.
- [26] T. L. Lai, "Information bounds and quick detection of parameter changes in stochastic systems," *IEEE Trans. Inf. Theory*, vol. 44, pp. 2917–2929, Nov. 1998.
- [27] A. G. Tartakovsky, M. Pollak, and A. Polunchenko, "Third-order asymptotic optimality of the generalized Shiryaev-Roberts changepoint detection procedures," *Theory of Prob and App.*, vol. 56, no. 3, pp. 534–565, 2011.